

COMP 105 Homework: Standard ML Modules

Due Wednesday, May 1 at 11:59 PM
(May be extended to May 2 by expending one token)

The purpose of this assignment is threefold:

- To get some practice with Standard ML modules
- To use Standard ML modules to put together a nontrivial program.
- To see how to reuse code that depends not just on other values, but also on other types.

You will complete problems 1–3 and problems A and B.

The code in this handout, together with a `compile` script, can be had by

```
git clone linux.cs.tufts.edu:/comp/105/git/ttt
```

The assignment is intended for **pairs**.

Contents

- [Part I: ML modules finger exercises](#)
 - [Part II: Playing adversary games](#)
 - ◆ [The idea behind the Abstract Game Solver \(AGS\)](#)
 - ◆ [Basic data in the problem: Players and outcomes](#)
 - ◆ [Specification of an abstract game](#)
 - ◆ [Adversary game problems](#)
 - ◇ [Problem A: Implement Tic-Tac-Toe](#)
 - ◇ [Testing your code: using the AGS with Tic-Tac-Toe](#)
 - ◇ [Problem B: Building an AGS](#)
 - [A common mistake to avoid when debugging your AGS](#)
 - ◇ [Using your AGS to play other games](#)
 - ◆ [Descriptions of the games](#)
 - ◇ [Tic Tac Toe](#)
 - ◇ [Nim](#)
 - ◇ [Connect 4](#)
 - ◆ [Extra Credit](#)
 - [What to submit](#)
 - [Acknowledgments](#)
 - [Appendix I: Two ways to compile Standard ML modules](#)
 - ◆ [Compiling Standard ML modules using Moscow ML](#)
 - ◆ [Compiling Standard ML to native machine code using MLton](#)
 - [Appendix II: The code we give you and how to get it](#)
 - [Appendix III: How your work will be evaluated](#)
 - ◆ [Interface design and modular programming](#)
 - ◆ [Mutability](#)
 - ◆ [Evaluation criteria](#)
 - ◇ [Interface design and documentation](#)
 - ◇ [Program structure](#)
 - ◇ [Performance and correctness](#)
-

Part I: ML modules finger exercises

1. **A simple functor (Difficulty *, Time **).** Remember the polymorphic sets we used when studying μ Scheme and core ML? To make them work, we tried two different strategies:

- ◆ To get set operations, pass in a comparison function and get a list of operations back ([Higher-order functions homework](#))
- ◆ Store the comparison function as part of the set itself ([Core ML homework](#))

The second strategy is not very reliable—imagine the chaos that ensues if you try to take the union of two sets, but they have *different* comparison functions!

The first strategy can be awkward in a typed language, because it can put extra type variables into the type of every function.

In this problem, you'll use Standard ML modules to define *dictionaries* which are guaranteed to use a consistent comparison function and do not require any extra parameters at run time. Instead you'll use one module parameter at link time.

In this problem, you'll be implementing dictionaries with this signature:

```
<dict-sig.sml>=
signature DICT = sig
  type key      (* a key used for lookup *)
  type 'a dict  (* a finite map from keys to values of type 'a *)

  exception NotFound of key

  val empty : 'a dict
  val find  : key * 'a dict -> 'a  (* may raise NotFound *)
  val bind  : key * 'a * 'a dict -> 'a dict

  (* contracts:
     empty is the empty map
     find (k, d) returns the x that d maps k to, or if d does not map k,
           it raises `NotFound k`
     bind (k, x, d) = d' such that
           - d' maps k to x
           - if k' <> k, d' maps k' the same way d does

  laws:
     find (k, bind (k, x, d)) = x
     find (k, bind (k', x, d)) = find (x, d)  if k <> k'
     find (k, empty)  raises NotFound k
  *)
end
```

Defines [bind](#), [DICT](#), [dict](#), [empty](#), [find](#), [key](#), [NotFound](#) (links are to index).

The comparison function is hidden inside the the module, so at run time, there are no extra types to think about and no comparison functions to pass!

To compile the file `dict-sig.sml` you will need to use the `-toplevel` option, e.g.,

```
mosmlc -toplevel -c dict-sig.sml
```

This problem has four parts:

- a. **Choose a representation** for dictionaries. Representations already implemented for you include the representation of environments in the book and the binary-search trees from the [Core ML](#) homework. Or you can come up with a new representation of your own.

COMP 105 Homework: Standard ML Modules

The abstraction is immutable, but if you want to use a mutable representation like a hash table, you can—but you must be sure that the mutability cannot be observed.

- b. Figure out **what you need to know about the key type** to implement your chosen representation. You probably will want one or more functions which consume values of type key and return values of type order. (Type order is explained in more detail in the Core ML homework.)
- c. **Define a signature** KEY to record the knowledge you figured out in the previous step.

Write this signature in file `key-sig.sml`.

- d. In file `dict.sml`, **implement a functor** `DictFn` which takes as argument a structure `Key` matching signature KEY and returns a structure matching DICT where type key = `Key.key`. **The dict type must remain abstract.**

To compile your functor, you will need to use the `-toplevel` option, and you'll need to provide compiled versions of the signatures for DICT and KEY.

```
mosmlc -toplevel -c dict-sig.ui key-sig.ui dict.sml
```

You will have produced `dict-sig.ui` by compiling `dict-sig.sml`, and similarly for `key-sig.ui`.

If you want to test your functor, define a structure `StringDict` `> DICT` where `type key = string`.

2. **One interface, three implementations (Difficulty *, Time **).** The ERROR signature from class looks like this:

```
<error-sig.sml>=
signature ERROR = sig
  type error      (* a single error *)
  type summary    (* summary of what errors occurred *)

  val nothing : summary                (* no errors occurred *)
  val oneError : error -> summary

  val <+> : summary * summary -> summary  (* combine summaries *)

  (* laws:
     nothing <+> s == s
     s <+> nothing == s
     s1 <+> (s2 <+> s3) == (s1 <+> s2) <+> s3    // associativity
  *)
end
```

Defines <+>, ERROR, error, nothing, oneError, summary (links are to index).

In this problem you will write three implementations of this signature:

- a. Implement a module

```
FirstError :> ERROR where type error = string
                    and type summary = string option
```

such that the summary always contains the *first* error to occur, if any.

Please put your implementation in file `first-error.sml`.

(My solution to this problem requires a dozen lines of code.)

- b. Implement a module

```
WorstError :> ERROR where type error   = { severity : int, msg : string }
                    and type summary = { severity : int, msg : string } option
```

such that the summary contains the *most severe* error to occur, if any.

Please put your implementation in file `worst-error.sml`.

(My solution to this problem requires 13 lines of code.)

- c. Implement a module

```
AllErrors :> ...
```

COMP 105 Homework: Standard ML Modules

such that the summary contains *all* the errors that occur. Please also **write and deploy a signature** suitable for sealing this module.

Please put your implementation in file `all-errors.sml`.

(My solution to this problem requires 11 lines of code.)

3. **Data structures (Difficulty *, Time **)**. A *heap* is a collection of *elements* with an operation that quickly finds and removes a minimal element. (The heap assumes that a total order exists on the elements; a minimal element is an element of the heap that is at least as small as any other element. A heap may contain more than one minimal element, in which case it is not specified which such element is removed.)

- a. *Design an abstraction* for representing heaps in ML. Your abstraction may be mutable or immutable.

Formalize your abstraction by giving an ML signature `HEAP` describing the abstraction. Be sure to

◇ Define two **abstract** types: one to represent a heap and one to represent an element.

◇ Identify each operation as a creator, producer, mutator, or observer (as described in Section 3.4.2 on page 76 of Ramsey)

◇ Specify what each operation does, either using informal English, algebraic laws, or both

The heap assumes that values of the element type are totally ordered; be sure to **expose that total order in the interface**.

Put your signature into a file called `heap-sig.sml`. You will need to compile it with the `-toplevel` option, e.g.,

```
mosmlc -toplevel -c heap-sig.sml
```

Notice that for this part of the problem you write **no code**. All you write is the interface.

- b. *Use your abstraction* to implement heap sort. That is, write a functor `HeapsortFn` that takes a structure matching signature `HEAP` and produces a structure that contains a function that sorts a list of elements by inserting all the elements into a heap, then removing them one by one until the heap is empty.

◇ Give your functor an explicit result signature, paying careful attention to type revelation.

◇ *You need not implement `HEAP`*. This is the whole point!

Put your `HeapsortFn` functor into a file called `heapsort.sml`. You will need to compile it with the `-toplevel` option, e.g.,

```
mosmlc -toplevel -c heap-sig.ui heapsort.sml
```

Because the `heapsort.sml` refers to signature `HEAP`, you must pass it the `heap-sig.ui` file where signature `HEAP` is defined. You will have produced `heap-sig.ui` by compiling `heap-sig.sml`.

Part II: Playing adversary games

In problems A and B below, you will implement and use a system for playing simple adversary games. The program will show game configurations, accept moves from the user and choose the best move.

The system is based on an abstract game solver (AGS) which, given a description of the rules of the game, will be able to select the best move in a particular configuration. An AGS is obtained by abstracting (separating) the details of a particular game from the details of the solving procedure. The solving procedure uses exhaustive search: it tries all possible moves and picks the best. Such a search can solve games of complete information, provided the configuration space is small enough. And the search is general enough that we can abstract away details of many games, separating the implementation of the solver from the implementation of the game itself.

To separate game from solver, in such a way that a single solver can be used with many games, requires a carefully designed interface. In this problem, we give you such an interface, which is specified using the SML signature `GAME`. (The signature was designed by George Necula and modified by Norman Ramsey.)

The `GAME` signature declares all the types and functions that an Abstract Game Solver must know about a game. The signature is general enough to cover a variety of games. Even details like "the players take turns" are considered to be part of

COMP 105 Homework: Standard ML Modules

the rules of the game—such rules are hidden behind the GAME interface, and the AGS operates correctly no matter what order players move in. (You could even implement a solitaire as a “two-player” game in which the second player never gets a turn!)

You will use two-player games in the last two parts of this assignment: implement a particular game and implement an AGS of your own.

The idea behind the Abstract Game Solver (AGS)

As players move, the state of a game moves from one *configuration* to another. In any given configuration, our solver considers all possible moves. After each move, it examines the resulting configuration and tries all possible moves from that configuration, and so on. In each configuration, the solver assumes that the player plays perfectly, that is, whenever possible the player will choose a move that forces a win.

This method (“exhaustive search”) is suitable only for very small games. Nobody would use it for a game like chess, for example. Nevertheless, variations of this idea are used successfully even for chess; the idea is to stop or “prune” the search before it goes too far.

Basic data in the problem: Players and outcomes

Representation is the essence of programming. We start by describing basic representations for the essential facts we assume about each game:

4. There are two *players*.
5. A game ends in an *outcome*: either one of the players has won, or the outcome is a tie.

The representations of these central concepts are *exposed*, not abstract. They are given by the signature PLAYER.

<player-sig.sml>=

```
signature PLAYER = sig
  datatype player = X | O      (* 2 players called X and O *)
  datatype outcome = WINS of player | TIE

  (* Returns the other player *)
  val otherplayer : player -> player
  val toString    : player -> string

  val outcomeToString : outcome -> string
end
```

Defines otherplayer, outcome, outcomeToString, PLAYER, player, toString (links are to index).

The signature player also includes some functions that compute with players and outcomes. Here's the implementation of signature PLAYER in a structure called Player.

<player.sml>=

```
structure Player :> PLAYER = struct
  datatype player = X | O
  datatype outcome = WINS of player | TIE

  fun otherplayer X = O
    | otherplayer O = X

  fun toString X = "X"
    | toString O = "O"

  fun outcomeToString TIE = "Tie"
    | outcomeToString (WINS p) = toString p ^ " wins"
```

COMP 105 Homework: Standard ML Modules

end

Defines otherplayer, outcome, outcomeToString, Player, player, toString (links are to index).

Although it might seem overly pedantic, we prefer to isolate details like the player names and how to convert them to a printable representation. To refer to Player types, constructors, and functions, you will use the "fully qualified" ML module syntax, as in the examples Player.otherplayer p, Player.X, Player.O, and Player.WINS p. The last three expressions can also be used as patterns.

Specification of an abstract game

The AGS can play any game that meets the specification given in signature GAME. This signature gives a contract for an entire module, which subsumes the contracts for all its exported functions.

<game-sig.sml>=

```
signature GAME = sig
  structure Move : sig (* information related to moves *)
    eqtype move          (* A move (perhaps a set of coordinates) *)
    exception Move      (* Raised (by makemove & fromString) for invalid moves *)
    val fromString : string -> move
                        (* converts a string to a move; If the string does not
                           correspond to a valid move, fromString raises Move *)
    val prompt : Player.player -> string
                (* Given a player, return a request for a move
                   for that player *)
    val toString : Player.player -> move -> string
                (* Returns a short message describing a
                   move. Example: "Player X moves to ...".
                   The message may not contain a newline. *)
  end

  type config          (* A representation for a game configuration. It
                        must include a full description of the state
                        of a game at a particular moment, including
                        keeping track of whose turn it is to move.
                        Configurations must appear immutable.
                        If a mutable representation is used, it must
                        be impossible for a client to tell that a
                        mutation has taken place. *)

  val toString : config -> string
                (* Returns an ASCII representation of the
                   configuration. The string must show whose turn it is. *)

  val initial : Player.player -> config
                (* Initial configuration for a game when
                   "player" is the one to start. We need the
                   parameter because the configuration includes
                   the player to move. *)

  val whoseturn : config -> Player.player
                (* Extracts the player whose turn is to move
                   from a configuration. We need this function because
                   the solver may need to know whose
                   turn it is, and the solver does not have
                   access to the representation of a configuration.
                   *)

  val makemove : config -> Move.move -> config
                (* Changes the configuration by making a move.
                   The player making the move is encoded in the
                   configuration. Be sure that the new
                   configuration knows who is to move. *)
```

COMP 105 Homework: Standard ML Modules

```
val outcome : config -> Player.outcome option
    (* If the configuration represents a finished game,
       return SOME applied to the outcome.
       If the game isn't over, return NONE. *)

val finished : config -> bool
    (* True if the configuration is final. This
       might be because one player has won,
       or it might be that nobody can move
       (which would be considered a tie). *)

val possmoves : config -> Move.move list
    (* A list of possible moves in a given
       configuration. ONLY final configurations
       might return nil. This means that a
       configuration which is not final MUST have
       some possible moves. In other words,
       part of the contract is that if 'finished cfg'
       is false, 'possmoves cfg' must return non-nil. *)

end
```

Defines config, finished, GAME, initial, makemove, Move, outcome, possmoves, toString, whoseturn (links are to index).

This is a broad interface. For example, there are three different ways to tell if a game is over!

Adversary game problems

Problem A: Implement Tic-Tac-Toe

A. Implement ``Tic-Tac-Toe.'' (Difficulty **, Time *)** More precisely, implement a module TTT matching signature GAME that describes Tic-Tac-Toe. If you are unfamiliar with Tic-Tac-Toe (elsewhere called ``Noughts and Crosses"), you can find an explanation at the [end of this assignment](#). Call your structure TTT, put it in the file `ttt.sml`, and use the following pattern :

```
<template for ttt.sml>=
structure TTT :> GAME =
  struct
    structure Move = struct
      type move = ... (* or use a datatype *)
      exception Move
      ...
    end

    type config = ... (* or use a datatype config = *)

    fun initial p = ...
    fun whoseturn c = ...

    ... and so on for all the values in GAME ...
  end
```

Defines TTT (links are to index).

Note the use of `:>`, which means that the *only* access to the types is through the functions in the GAME signature.

When writing TTT, you must define *all* types and values mentioned in the signature GAME, and all values must have the types specified. You might want to define additional values, which you will be able to use as helper functions. These functions cannot be called from anyone else's code: because TTT is forced to have signature GAME, the functions are not visible outside

COMP 105 Homework: Standard ML Modules

the `TTT` module, and therefore no other code can depend on them.

So we can test your code, we insist that you use the following names of squares in `Move.toString` and `Move.fromString`:

```
upper left | upper middle | upper right
-----+-----+-----
middle left | middle | middle right
-----+-----+-----
lower left | lower middle | lower right
```

You should always print and recognize these full names. If you wish, you may also recognize the abbreviations `ul`, `um`, `ur`, `ml`, `m`, `mr`, `ll`, `lm`, and `lr` in the function `Move.fromString`.

Here are step-by-step instructions:

- Choose how you will represent the state of the game (i.e., define `config`). This step is crucial because it determines how complex your implementation will be. There are many possible representations; any one is OK provided you are able to implement the functions required by the signature. Choose a representation that will make it easy to implement `makemove`, `possmoves`, and `outcome`.

The AGS cannot possibly depend on your choice of representation (the ML module system guarantees it), so you are free to choose whatever representation you like. Even more important, **you can change your representation at any time**, and no code outside your own module will be affected. If you have any difficulty implementing the functions in the `GAME` interface, you *should* change your representation—or at least think about it.

Document your representation by stating any invariants that it satisfies, and explain how your representation relates to the abstraction of the tic-tac-toe grid.

You might be tempted to use mutable data to represent game state. **Don't!** The contract of the `GAME` interface requires that any value of type `config` be available to the AGS indefinitely. Mutating a configuration is not safe.

If you think you might want *immutable* arrays, check out the `Vector` structure (see the [ML supplement](#)). (You can find out what's in any ML structure by typing, e.g., `open Vector` at the interactive prompt, or you can consult the [Standard Basis documentation](#). You can also use Moscow ML's help system, e.g,

```
- help "Vector";
```

If you get interested in vectors, don't overlook the function `Vector.tabulate`.)

One more thing. You may be tempted to start out by representing the contents of a square on the board using 0 and 1 or other arbitrary values. If you go this route, why not use `Player.player` option? It will make your program more elegant and easier to understand.

- Choose a representation for moves. That is, write `move`. Everything said for configurations applies here also, but this choice seems less critical.
- Declare the exception `Move`.
- Write the function `initial`.
- Write the function `whoseturn`.
- Write `makemove`. The contract requires it to be Curried.
- Write `outcome`. If the configuration is not final and nobody has won, return `NONE`.

Hints for Tic-Tac-Toe:

- ◆ You could write a function which checks lines, another that checks columns and finally one that checks diagonals. Then `outcome` could call these functions with the right parameters.
 - ◆ You could try pattern matching. Standard ML supports pattern matches on vectors by, e.g., `case a of #[x, y, z] => ...`
- Write `finished`. This function should return true if somebody has won or if no move is possible (everybody is stuck). Be smart and use another function to do most of the work.

COMP 105 Homework: Standard ML Modules

- i. Write `possmoves`. This function must return a list of the possible moves (in no particular order). It is in everybody's interest that the list have no duplicates. *If the game is over, no further moves are possible*, and `possmoves` must return `nil`. (In this case, according to contract, `finished` must return `true`.)

If you want to be clever, you can exploit rotation and reflection symmetries to prune the list returned by `possmoves`. You may be surprised how much difference this makes to performance. **For extra credit**,

1. submit a version of `possmoves` that exploits symmetry to minimize the number of possible moves
2. give a "back of the envelope" estimate of the time to be saved when the AGS plays against itself
3. measure the actual time savings using the `Timer` and `Time` structures thusly:

```
fun time f arg =
  let val start = Timer.startRealTimer()
      val answer = f arg
      val endit = Timer.checkRealTimer start
  in print ("Time is " ^ Time.toString endit ^ "\n");
    answer
  end
```

You can also try `startCPUTimer` and `checkCPUTimer`, but the answers you get are a bit more complicated.

- j. Create a data structure that associates each move with one or more representations of that move as strings. This data structure will be a single point of truth, and it will ensure that `Move.fromString` and `Move.toString` cannot possibly be inconsistent, even if you make a mistake in their implementations. **To get full credit for this problem**, your code must provide this guarantee by ensuring there is a single point of truth.
- k. Using your data structure from the previous step, write `Move.toString`. This function must return a string of the form "Player... moves to ..." which does *not* end in a newline. You can build your strings using concatenation (^) and exported functions from other modules (e.g. `Player.toString`). To convert integer values to strings you can use the function `Int.toString`.
- l. Write `toString`. You must return a simple ASCII representation of the state of the game configuration. The value should end in a newline. Don't forget to include the player whose turn it is to move. Give us more than a simple list of numbers. You can print a nice little "ASCII graphics" layout using only a few characters. To get you started, here is some untested sample code to print a row; it has type `player option list -> string`:

```
<sample function rowString>=
local
  fun boxString (SOME p) = Player.toString p
    | boxString (NONE ) = " "
in
  fun rowString [] = "\n"
    | rowString (box :: boxes) = "|" ^ boxString box ^ " " ^ rowString boxes
end
```

Defines `boxString`, `rowString` (links are to index).

`Move.toString` and `toString` are not involved in the correctness of the AGS; they are used by the interactive player to show you what's happening. The better your output, the more fun it will be to play. You can see a simple sample by running `/comp/105/bin/ttt`.

- m. Write `Move.prompt`. It takes the player whose turn it is to move, and it returns a prompt message (without newline) asking the specified player to give a move in the format we specified (naming the square).
- n. Write `Move.fromString`. This function should take a string (which is probably the reply given after a call to `Move.prompt`), and it should return the move corresponding to that string. If there is no such move, it should raise an exception.

You should try to write `Move.fromString` in such a way that `Move.fromString` and `Move.toString` cannot possibly be inconsistent, even if you make a mistake. Because we give you a lot of freedom, it is hard to specify precisely what it means to be consistent, but here is a rough specification:

COMP 105 Homework: Standard ML Modules

For any move `m`, there should be an `i` and `n` such that `m` is equal to `Move.fromString (String.extract (Move.toString m, i, SOME n))`.

Be sure to try your functions on simple configurations.

Hints: You may find it useful to define a structure `Grid` that you can use to represent a square or rectangular array of values of type 'a. Defining suitable analogs of `map` and `fold` on the grid will help, as will functions to extract sub-grids (rows and columns). If you then define reflection and rotation on grids, you can easily do the extra credit.

The most common mistake on this problem is to permit players to continue to move even when the game is over.

Bob Harper's code for Tic-Tac-Toe is 146 lines of Standard ML. I have a slicker version at only 87 lines—and it is four times faster. It works by exploiting bit-level parallelism using the `Word` structure and by flagrantly disregarding most of the hints given above.

Testing your code: using the AGS with Tic-Tac-Toe

To build a version of the AGS for ``Tic-Tac-Toe'' you must use the following command:

<example of creating a game-specific AGS>=

```
structure TTTAgs = AgsFun(structure Game = TTT)
```

Defines `TTTAgs` (links are to index).

Of course, I can't do any of this until I use the Moscow ML load function to get access to `AgsFun` and `TTT`. Here is an example:

<transcript from an actual session>= [D->]

```
: nr@labrador 7147 ; mosml
Moscow ML version 2.10-2 (Tufts University, February 2011)
Enter `quit();' to quit.
- load "ags";
> val it = () : unit
- load "ttt";
> val it = () : unit
- structure TTTAgs = AgsFun(structure Game = TTT);
> structure TTTAgs :
  {structure Game :
    {structure Move :
      {type move = move,
        exn Move : exn,
        val fromString : string -> move,
        val prompt : player -> string,
        val toString : player -> move -> string},
      type config = config,
      val finished : config -> bool,
      val initial : player -> config,
      val makemove : config -> move -> config,
      val outcome : config -> outcome option,
      val possmoves : config -> move list,
      val toString : config -> string,
      val whoseturn : config -> player},
    val bestmove : config -> move option,
    val forecast : config -> outcome}
  }
-
```

This functor application creates a structure that implements the `AGS` signature:

<ags-sig.sml>=

```
signature AGS = sig
  structure Game : GAME
  (* Given a configuration, returns the
```

COMP 105 Homework: Standard ML Modules

```
      * most beneficial move for the player
      * to move *)
val bestmove : Game.config -> Game.Move.move option

      (* Given a configuration, returns the
      * best possible outcome for the player
      * whose turn it is, assuming opponent
      * plays optimally *)
val forecast : Game.config -> Player.outcome
end
```

Defines [AGS](#), [bestmove](#), [forecast](#), [Game](#) (links are to index).

The function [bestmove](#) returns the best move in a configuration, or NONE if no move is possible, i.e., the configuration is final. The function [forecast](#) predicts the outcome from a configuration if both players make perfect moves.

These functions can be *slow* because the AGS tries all possible combinations of moves. Be patient.

We have also provided you an interactive player. It uses the AGS so you must instantiate it to the Tic-Tac-Toe AGS using the following command:

```
<examples>= [D->]
structure P = PlayFun(structure Ags = TTAgs);
```

Defines [P](#) (links are to index).

Again, to get PlayFun you will have to load the right module:

```
<transcript from an actual session>+= [C-D]
- load "play";
> val it = () : unit
- structure P = PlayFun(structure Ags = TTAgs);
> structure P :
  {structure Game : ...
    exn Quit : exn,
    val getamove : Player.player list -> Game.config -> move,
    val play : (config -> move) -> config -> outcome}
-
```

The structure this application creates implements the following signature :

```
<play-sig.sml>=
signature PLAY = sig
  structure Game : GAME
  exception Quit
  val getamove : Player.player list -> Game.config -> Game.Move.move
    (* raises Quit if human player refuses to provide a move *)

  val play : (Game.config -> Game.Move.move) -> Game.config -> Player.outcome
end
```

Defines [Game](#), [getamove](#), [PLAY](#), [play](#), [Quit](#) (links are to index).

The function [getamove](#) expects a list of players for which the computer is supposed to play (the computer might play for X, for O, for both or for none). The return value is a function which the interactive player will use to request a move given a configuration. The idea is that the function returned will ask the AGS for a move if the computer is playing for the player to move, or will prompt the user and convert the user's response into a move.

The function [play](#) expects an input function (one built by [getamove](#)) and a starting configuration. This function then starts an interactive loop printing the intermediate configurations and prompting the users for moves (or asking the AGS where appropriate). One example is :

COMP 105 Homework: Standard ML Modules

```
<examples>+= [←D]  
val computerxo = P.getamove [Player.X, Player.O]  
    (*Computer plays for both X and O *)  
  
val computero = P.getamove [Player.O]  
    (*Computer plays only O *)  
  
val cnfi = TTT.initial Player.X  
    (* Empty configuration with X to start *)  
  
val frustration = P.play computero  
    (* We play against the computer *)  
  
val _ = frustration cnfi  
    (* A frustrating exercise *)
```

Defines [cnfi](#), [computero](#), [computerxo](#), [frustration](#) (links are to index).

Problem B: Building an AGS

B.Implement an Abstract Game Solver (Difficulty *).** Given a configuration, an AGS should compute the *benefits* of all possible moves and pick the best one. More precisely, given a configuration and a player, the AGS assigns a benefit to that player of that configuration. A final configuration in which X has won should have maximum benefit to X and minimum benefit to O, and vice versa. Ties should have intermediate and equal benefit to both players. We compute the benefit of an intermediate configuration by looking at all possible moves and the benefits of the resulting configurations.

There are a variety of ways to view benefits; for example, we could assign larger benefits to winning quickly, and so on. For this assignment, however, it will be sufficient to consider three levels of benefits:

- Player to move can force a win
- Both players can force a tie
- Player to move can be forced to lose by his adversary

For **extra credit** you can prove that one of these three situations must hold in any game described by the [GAME](#) signature, provided that the game is deterministic and is guaranteed to terminate after finitely many moves.

Write an AGS using the following template:

```
<template for functor AgsFun>=  
functor AgsFun (structure Game : GAME) :> AGS  
  where type Game.Move.move = Game.Move.move  
    and type Game.config = Game.config  
= struct  
  structure Game = Game  
  
  fun bestresult conf = ...  
  fun bestmove conf = ...  
  fun forecast conf = ...  
end
```

Defines [AgsFun](#), [bestmove](#), [bestresult](#), [forecast](#), [Game](#) (links are to index).

Note how annoying the where type declarations are: they look tautological, but they're not. Complain to Dave MacQueen and Bob Harper.

We recommend you create a helper function [bestresult](#) with type

```
val bestresult : Game.config -> Game.move option * result
```

COMP 105 Homework: Standard ML Modules

where `result` is a representation you choose.

The idea is that `bestresult conf = (bestmove, whathappens)` where

- If the player can't move, `bestmove` is `NONE`.
- If the player can move, `bestmove` is `SOME m`, where `m` is the best possible `Game.move` for the player in this configuration.
- Value `whathappens` explains what the AGS predicts is the outcome of the game if both players play perfectly. It suffices to use a result of type `Player.outcome`, but you can play around with this one some—for example, you might want to return an outcome like ```Player X wins in 3 moves.`" This would help you build an aggressive AGS.

You might be tempted to use a ```relative`" outcome like ```Win, Lose, or Tie.`" This can be made to work, but it is harder to get right, especially in games where players don't always take turns.

In order to make `bestresult` work, you'll need some recursive calls. You'll also want a helper function that lets you compare the benefits of different outcomes, so `bestresult` can choose the most desirable outcome for the current player.

Hints:

- To speed up the AGS, you may want to stop the search as soon as you find a forced win.
- Do **not** assume that players take turns, that the last player to move always wins, or any other properties of Tic-Tac-Toe. Use `whoseturn` and `outcome` instead. We will test your AGS on games that are quite different from Tic-Tac-Toe.

To test your AGS, you'll need to replace our `ags.ui` and `ags.uo` files with the ones you compile from your source code. At this point you'll be able to run the same test cases you used earlier, as well as what's in part B.

My AGS takes about 40 lines of Standard ML.

A common mistake to avoid when debugging your AGS

If you build a simple AGS that fits in 40 lines of code, it is not going to try to fool you: if you can force a win, the AGS will pick a move more or less arbitrarily. A simple AGS has no notion of ```better`" or ```worse`" moves; it knows only whether it can force a win.

Here's the common mistake: you're playing against the AGS, and it makes a terrible move. You think it's broken. For example, suppose you are playing X, the AGS is playing O, and you start play in this position:

```
-----  
|   | O |   |  
-----  
|   | X |   |  
-----  
|   |   |   |  
-----
```

You move in the upper left corner. **The AGS does not move lower right to block you.** Is it broken? No—the AGS recognizes that you can force a win, and it just gives up.

If you want an AGS that won't give up, for extra credit you can implement an aggressive version that will delay the inevitable as long as possible. An aggressive AGS searches more states so that it can (a) win as quickly as possible and (b) hold on in a lost position as long as possible.

My aggressive AGS is under 60 lines of Standard ML code.

Using your AGS to play other games

The code we supply includes a description of the game ``Nim". The structure that implements ``Nim" is called `structure Nim`. After you create an AGS solver and an interactive player for ``Nim" you can play Nim with the AGS. The commands to instantiate AGS to ``Nim" are:

`<nim examples>=`

```
structure NIMAgS = AgsFun(structure Game = Nim)
structure PN = PlayFun(structure Ags = NIMAgS)
```

Defines `NIMAgS`, `PN` (links are to index).

You play Nim by running `/comp/105/bin/nim`, but the user interface stinks.

We've also implemented a version of ``Connect 4" that would be better called ``Connect 3" (since 4 would be too slow). It is in `/comp/105/bin/four`.

Descriptions of the games

Here are descriptions of 3 games: ``Tic-Tac-Toe", ``Nim" and ``Connect 4". Do not worry if you haven't seen these games before—you can learn by playing against a perfect or near-perfect player. (The Connect 4 player would be perfect if it were faster.) For the purpose of this assignment you do not have to know any tricks of the games but only to understand their rules.

Tic Tac Toe

This is an adversary game played by two persons using a 3x3 square board. The players (traditionally called X and O) take turns in placing X's or O's in the empty squares on the board (player X places only X's and O only O's). The board is empty in the initial configuration.

The first player who managed to obtain a full line, column or diagonal marked with his name is the winner. The game can also end in a tie. In the picture below the first configuration is a win for O, the next two are wins for X and the last one is a tie.

```
-----
| X |  | X |  |  |  |  |  |
-----
|  | X |  |  |  |  |  |  |
-----
| O | O | O |  |  |  |  |
-----

-----
|  |  | X |  |  |  |  |  |
-----
| O | X | O |  |  |  |  |
-----
| X | O |  |  |  |  |  |
-----
| X |  | O |  |  |  |  |
-----

-----
| O | O | X |  |  |  |  |
-----
| X | X | O |  |  |  |  |
-----
| O | X | O |  |  |  |  |
-----
```

In this game a player who plays perfectly cannot lose. All your base are belong to the AGS.

Nim

This is an adversary game played by two persons. The game is played with number of sticks arranged in 3 rows. In the initial state the rows usually contain 3, 5 and 7 sticks respectively. The players take turns in removing sticks: each player can remove 1, 2 or 3 adjacent sticks from one row. The one that removes the last stick is the loser. Or, stated differently the first player who has no sticks to remove is the winner. Below were presented two configurations. The first one is the initial configuration (for the 3, 5 and 7) case and the other one is the configuration obtained after a few moves. A possible sequence of moves that might lead to this configuration is:

1. X removes sticks 0, 1 and 2 from row 1
2. O removes stick 1 from row 0
3. X removes stick 6 from row 2
4. O removes sticks 3 and 4 from row 2

COMP 105 Homework: Standard ML Modules

```
Row 0: | | |           | _ |
Row 1: | | | | |     _ _ _ | |
Row 2: | | | | | | | | | | _ _ | _
```

We have represented a stick using a ``|'' and a missing stick using a ``_'. It might be wise to play with a smaller configuration (2, 3 and 4 for example) because otherwise the AGS will take too long to produce its answers.

For this game the first player can always win no matter what the other does. If you let the AGS start you have no chance. If you play first you can beat the AGS, but you have to play well.

Connect 4

This is an adversary game played by two persons using 6 rods and 36 balls. Imagine the rods standing vertically, and each ball has a hole in it, so you can drop a ball onto a rod. The balls are divided in two equal groups marked X and O. The players take turns in making moves. A move for a player consists in sliding one of its own balls down a rod which is not full (the capacity of a rod is 6). The purpose is to obtain 4 balls of the same type adjacent on a horizontal, vertical or diagonal line. The game ends in a tie when all the rods are full and no player has won. We represent below the initial configuration of the game and a final state where X has won.

```
| | | | | |   | | | | | |
| | | | | |   | | | | | |
| | | | | |   | | | | | |
| | | | | |   O | | | | |
| | | | | |   O | O | | |
| | | | | |   O X X X X |
-----
```

Our version uses 5 rods and connects 3, because otherwise the AGS takes too long.

Extra Credit

Symmetry. Speed up Tic-Tac-Toe by exploiting symmetry as described [above](#).

Proof. Prove the ``forcing'' property of these simple games as described above [above](#).

Four. Implement Connect 4.

Game. Suggest another simple adversary game, and (with the instructor's approval) implement it. The game should be small with a small number of possible moves; otherwise the exhaustive search is infeasible.

Aggression. With the simple benefits outlined above, the AGS will ``give up'' if it can't beat a perfect player---all moves are equally bad, and it apparently moves at random. What this scheme doesn't account for is that the other player might not be perfect, so there is a reason to prefer the most distant loss. In the dual situation, when the AGS knows it can win no matter what, it will pick a winning move at random instead of winning as quickly as possible. **This behavior may lead you to suspect bugs in your AGS. Don't be fooled.**

Change your benefits so that the AGS prefers the closest win and the most distant loss. (This means you can only prune the search if you find a win in one move.) If you are clever, you can encode all this information in one value of type `real`.

Learning. We can re-use the `GAME` signature for more than one purpose. Implement a ``matchbox'' learning engine in the style explained by [Martin Gardner's article](#) on the reading list. You can use the SML/NJ library to store state with each configuration, using the following signature:

```
signature ORDERED_GAME = sig
  include ORD_KEY
  include GAME
```

COMP 105 Homework: Standard ML Modules

```
sharing type conf = ord_key
end;
```

You may have to modify the AGS to notify each player of the outcome of the game. See me for more help with details.

What to submit

For this assignment you should use the script `submit105-sml` to submit

- Either `README`
- For problem 1, files `key-sig.sml` and `dict.sml`
- For problem 2, files `first-error.sml`, `worst-error.sml`, and `all-errors.sml`
- For problem 3, files `heap-sig.sml` and `heapsort.sml`
- For problem A, file `ttt.sml`
- For problem B, file `ags.sml`

The ML files should contain all structure and function definitions that *you* write for this assignment (including any helper functions that may be necessary), in the order they should be compiled. The files you submit must compile with Moscow ML, using the compile script we give you. We will reject files with syntax or type errors. Your files should compile *without warning messages*. If you must, you can include multiple structures in your files, but *please don't make copies of the structures and signatures above*; we already have them.

Acknowledgments

This assignment is derived from one graciously provided by [Bob Harper](#). [George Necula](#), who was his teaching assistant at the time (and is now a professor at Berkeley and is world famous as the inventor of proof-carrying code), did the bulk of the work.

Appendix I: Two ways to compile Standard ML modules

Compiling Standard ML modules using Moscow ML

To compile an individual module using Moscow ML, you type

```
mosmlc -c -toplevel filename.sml
```

This puts compiler-interface information into `filename.ui` and implementation information into `filename.uo`. Perhaps surprisingly, either a signature or a structure will produce *both* `.ui` and `.uo` files. This behavior is an artifact of the way Moscow ML works; it should not alarm you.

If your module depends on another module, you will have to mention the `.ui` file on the command line as you compile. For example your `DictFn` functor depends on both `DICT` and `KEY` signatures. If `DictFn` is defined in `dict.sml`, `KEY` is defined in `key-sig.sml`, and `DICT` is defined in `dict-sig.sml`, then to compile `DictFn` you run

```
mosmlc -toplevel -c dict-sig.ui key-sig.ui dict.sml
```

To talk about what happens after you compile, I'll use another example:

```
mosmlc -c -toplevel game-sig.ui player.ui ttt.sml
```

This compilation produces two files:

- `ttt.ui`, which can be used on the command line when compiling other units that depend on `TTT`.
- `ttt.uo`, which contains the compiled binary

COMP 105 Homework: Standard ML Modules

You can do two things with the `.uo` files:

- When you are debugging, it is tremendously useful to get compiled modules into the interactive system. Load them directly using `load`, e.g.,

```
: nr@labrador 2856 ; mosml
Moscow ML version 2.10-2 (Tufts University, February 2011)
Enter `quit();' to quit.
- load "ttt";
> val it = () : unit
- open TTT;
> structure Move :
  {type move = move,
   exn Move = Move : exn,
   ...}
type config = config
...
val finished = fn : config -> bool
...
```

Once you load a module, you cannot recompile it and reload it later. Loading it again has no effect, even if the code has changed; you have to start Moscow ML over again.

- You can use `mosmlc` to link a bunch of `.uo` files together to form an executable binary. To do anything interesting, one of the source files should have a top-level call to `play`, `forecast`, or some other interesting function.

Here is an example of a command line I use on my system to build an interactive game player:

```
mosmlc -toplevel -o games player-sig.uo player.uo game-sig.uo \
      ags-sig.uo play-sig.uo slickttt.uo ttt.uo \
      ags.uo aggress.uo nim.uo four.uo peg.uo mrun.uo
```

Order matters; for example, I have to put `player.uo` *after* `player-sig.uo` because the `Player` structure defined in `player.sml` uses the `PLAYER` signature defined in `player-sig.sml`.

The `git` repository for this assignment includes a `compile` script that may help with compiling Moscow ML modules.

Compiling Standard ML to native machine code using MLton

If your games are running too slow, compile them with `MLton`. `MLton` is a whole-program compiler that produces optimized native code. To use `MLton`, you list all your modules in an MLB file, and `MLton` compiles them at one go. We provide a `ttt.mlb` file, so you can compile with, e.g.,

```
mlton -output ttt -verbose 1 ttt.mlb
```

Because `MLton` requires source code, you will be able to use it only once you have your own AGS. More information about `MLton` is available on the man page and at www.mlton.org.

Appendix II: The code we give you and how to get it

In the `git` repository at `/comp/105/git/ttt`, you'll find sources for most of the signatures, structures, and functors in this assignment. You'll also find an AGS in binary form only. And you'll find a `compile` script, which should compile your code using the Moscow ML compiler, `mosmlc`, and a `ttt.mlb` file for `MLton`.

Appendix III: How your work will be evaluated

Interface design and modular programming

Modules provide **controlled information hiding**. To use the power of modules effectively, design interfaces that are both **usable** and **abstract**:

- Enough types should be made manifest that the module can be used.
- No more types should be made manifest than is absolutely necessary.

Here's an acid test:

- If an interface exports an abstract type like `heap` or `queue` or `dictionary` or `computation`, it must be possible to create a value of the abstract type using the functions exported by the interface. It must *not* be possible to create a value of the abstract type *without* using the functions exported by the interface.

Other things being equal, hiding more information is better. When in doubt, make a type abstract.

Mutability

An abstract type, in any language, is always either **mutable** or **immutable**. The choice is made by the designer of the interface, and the choice is available regardless of programming language.

- *Mutable* abstractions have "object identity": if you mutate x and you see the change when you look at y , then x and y are the same object. For example, if sets of numbers are mutable abstractions, then just because x and y both represent the empty set doesn't mean they are the same object.
- *Immutable* abstractions lack object identity: if the result of observing a property of x is always identical to the result of observing the same property of y , then they may or may not be the same object, but it doesn't matter, because no program can tell the difference between them. For example, "natural number" is an immutable abstraction, and if you have the natural number ~ 7 , it makes no sense to ask "which 7 do I have," because they're all indistinguishable.
- The interface for a mutable abstraction always provides at least one *creator* function which creates a new object that is distinct from any other. The interface for an immutable abstraction doesn't require creators. For example, if you implement mutable lists, you need a way to allocate a new, empty list. But if lists are *immutable*, you can just provide one empty-list value and forget about it.
- The interface for a mutable abstraction provides mutators that change the contents or value of an object. The return type of a mutator typically does *not* include a value of the abstract type; instead of a returning a new value, the mutator changes one of its arguments.

The interface for an immutable abstraction typically provides *producers* that take in an existing object and produce a new object, which is returned. The return type of a producer always includes a value of the abstract type.

- In a statically typed language like ML, it is usually not necessary to say explicitly whether an abstraction is mutable or immutable—it should be obvious from the types.

The key takeaway here is that any abstraction you create should be either mutable or immutable, and **the type of every operation must be consistent with your decision**.

Evaluation criteria

Interface design and documentation

We'll look most closely at the design and documentation of your interfaces.

	Exemplary	Satisfactory	Must improve
Interface design	• Signatures defined in the program make types abstract when possible but manifest when necessary.	• A signature makes a type abstract but provides so many operations with such detailed contracts that only one representation is possible,	• A signature makes a type manifest that should be abstract.

COMP 105 Homework: Standard ML Modules

	<ul style="list-style-type: none"> • In any signature, no type is made manifest unless it is absolutely necessary to expose the type. • Every signature that is used to seal a module exposes a combination of <i>operations</i> and <i>types</i> so that the abstractions described in the signature can be used. • If an interface exports an abstract type, it is possible to <i>create</i> a value of the type using the functions exported by the interface. • If an interface exports an abstract type, it is possible to <i>observe</i> (or consume or query) a value of the type using the functions exported by the interface. • Every abstract type is supported by enough functions that it is possible to observe any property that would be observable in the world of ideas. • Every signature that describes an argument to a functor is as small as possible, but no smaller. • In every signature, every abstract type is either clearly mutable or clearly immutable. • Every abstraction handles polymorphism either through the core language (using a type parameter) or through the module language (using a functor parameter), never a mix of both. 	<p>and the representation might as well be exposed.</p> <ul style="list-style-type: none"> • A type is exposed (made manifest) unnecessarily, but the value constructors are hidden, so an adversary cannot forge a value of the type. • Some properties of an abstract type are observable and some aren't. • A signature that describes an argument to a functor contains superfluous elements that are not needed by the functor. • A polymorphic abstraction uses a type parameter to achieve polymorphism, but the interface is cluttered because one or more operations on the abstract type have to take an extra parameter (for example, a comparison function). 	<ul style="list-style-type: none"> • A type is so exposed that an adversary can create a value of the type without using the functions in the interface, possibly violating the type's invariants. • An abstract type named in a signature can't be used because there are no operations available to create values of that type. • An interface provides functions to create a value of abstract type, but afterward the value can't be observed: there is nothing you can <i>do</i> with it. • A signature that describes an argument to a functor is missing elements that are required for the functor to do its job. • Some of the operations on an abstract type have types which suggest that the abstraction is mutable; others suggest it is immutable. • An interface appears to provide a polymorphic abstraction using a core-language type parameter, but closer inspection shows that the type parameter performs no useful function; the interface requires the use of an abstract type also named in the interface.
Interface documentation	<ul style="list-style-type: none"> • In HEAP, every operation is explicitly classified as a creator, producer, observer, or mutator. • Every signature documents what its abstract types stand for in the world of ideas. (Sometimes the name of the type is sufficient.) 	<ul style="list-style-type: none"> • Heap operations are grouped by classification, but the classification is not explicit. • A function raises an exception and it's not documented explicitly, but based on the name and type of the function and the names of exported exceptions, course staff can make a reliable guess. 	<ul style="list-style-type: none"> • Heap operations are not classified as creators, producers, observers, or mutators. • A function raises an exception that can't be caught, because the exception isn't exported in the interface.

COMP 105 Homework: Standard ML Modules

	<ul style="list-style-type: none"> • In every signature, the contract of every function is either implicit in its name and type, or is stated explicitly. (Explicit statements may refer to the world of ideas or may be composed by writing algebraic laws.) • If an exported function can raise an exception, this behavior is documented as part of the function's contract. • Exceptions that might be raised are either part of the interface or are defined in the context (usually the initial basis) 	
--	---	--

Program structure

We'll be looking for you to seal all your modules and to isolate the truth about Tic-tac-toe moves in just one point in your code. We'll also be looking for the usual hallmarks of good ML structure.

	Exemplary	Satisfactory	Must improve
Structure	<ul style="list-style-type: none"> • All modules are sealed using the opaque sealing operator <code>></code> • There is exactly one place in the Tic-tac-toe code where the internal representation of each move is related to its external string representation(s) • Code uses basis functions effectively, especially higher-order functions on list and vector types. • Code has no <u>redundant case analysis</u> • Code is no larger than is necessary to solve the problem. 	<ul style="list-style-type: none"> • Most modules are sealed using the opaque sealing operator <code>></code> • Internal representations and string representations of Tic-tac-toe moves are related in multiple places in the code, but all relations are immediately adjacent to their inverse relations. • Code uses the familiar functions, but misses opportunities to use unfamiliar functions like <code>Vector.tabulate</code>. • Code has one <u>redundant case analysis</u> • Code is somewhat larger than necessary to solve the problem. 	<ul style="list-style-type: none"> • Only some or no modules are sealed using the opaque sealing operator <code>></code> • A module is defined without ascribing any signature to it (unsealed) • Conversion between internal representations and string representations of Tic-tac-toe moves is done by separate functions which don't share code or data structures but are required to be inverses. Information is duplicated. • Code misses opportunities to use <code>map</code>, <code>fold</code>, or other familiar HOFs. • Code has more than one <u>redundant case analysis</u> • Code is almost twice as large as necessary to solve the problem. • <i>Or</i>, code contains near-duplicate functions (most likely in AGS)

Performance and correctness

Finally, we'll look to be sure your code meets specifications, and that the performance of your AGS is as good as reasonably possible.

COMP 105 Homework: Standard ML Modules

	Exemplary	Satisfactory	Must improve
Correctness	<ul style="list-style-type: none"> • The HEAP signature contains only operations that are appropriate to a set with a known minimum element. • The HEAP signature exposes the total order that exists on values of the element type. • Tic-tac-toe code fulfills the contracts specified in the GAME and MOVE signatures. • AGS code makes no additional assumptions about the implementations of Player, Move, or Game. 	<ul style="list-style-type: none"> • The HEAP signature contains superfluous operations. • Tic-tac-toe code fulfills the contracts specified in the GAME and MOVE signatures, except that it continues to offer moves even after the game is over. 	<ul style="list-style-type: none"> • The HEAP signature contains superfluous operations that the course staff believe are intended to make sorting easier. • The HEAP signature contains a sort function (serious fault). • The HEAP signature contains a higher-order function that visits heap elements in total order (serious fault). • The HEAP signature does not expose the total order that exists on values of the element type. • Tic-tac-toe code violates one of its contracts. • AGS code assumes that players take turns.
Performance	<ul style="list-style-type: none"> • The AGS implements its <code>bestmove</code> and <code>forecast</code> functions using a single, pruned search that stops once the best move or outcome is known. • <i>Or</i>, the AGS implements <code>bestmove</code> and <code>forecast</code> by making just <i>one</i> search through the state space of possible game configurations. 	<ul style="list-style-type: none"> • Function <code>bestmove</code> or <code>forecast</code> may search the state space of possible configurations more than once. 	<ul style="list-style-type: none"> • Function <code>bestmove</code> or <code>forecast</code> may search the state space of possible configurations more than twice.